REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302 and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
August 1996

3. REPORT TYPE AND DATES COVERED
Technical Report

4. TITLE AND SUBTITLE
Bayesian Knowledge-Bases

5. FUNDING NUMBERS

6. AUTHOR(S)
Eugene Santos Jr. and Eugene S. Santos

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Institute of Technology, WPAFB OH 45433-6583

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT/EN/TR96-05

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution Unlimited

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Abstract

Managing uncertainty in complex domains requires a flexible and semantically sound knowledge representation. This is especially important during the initial knowledge engineering and subsequent maintenance of the knowledge base. We present a new model of knowledge representation called Bayesian Knowledge Bases. It unifies a "if-then" style rules with probability theory. We can prove that such a merger remains fully probabilistic and yet maintains full flexibility and intuitiveness.

DTIC QUALITY INSPECTED 4

14. SUBJECT TERMS
Knowledge Representation, Uncertainty, Knowledge Organization, Knowledge Consistency, Probabilistic Representation, Expert Systems

15. NUMBER OF PAGES
23

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

UNCLASSIFIED

AFIT/EN/TR96-05

Air Force Institute of Technology

Bayesian Knowledge-Bases

Eugene Santos Jr.   Eugene S. Santos

August 12, 1996

Approved for public release; distribution unlimited

19970502 236

Bayesian Knowledge-Bases

Eugene Santos Jr.[1] and **Eugene S. Santos**

August 11, 1996

Abstract

Managing uncertainty in complex domains requires a flexible and semantically sound knowledge representation. This is especially important during the initial knowledge engineering and subsequent maintenance of the knowledge base. We present a new model of knowledge representation called Bayesian Knowledge Bases. It unifies a "if-then" style rules with probability theory. We can prove that such a merger remains fully probabilistic and yet maintains full flexibility and intuitiveness.

Keywords: Knowledge Representation, Uncertainty, Knowledge Organization, Knowledge Consistency, Probabilistic Representation, Expert Systems

[1] This research was supported in part by AFOSR Project #940006 and #9600989.

1 Introduction

Managing uncertainty in complex domains continues to remain a difficult task especially during knowledge acquisition and verification and validation. There are a wide variety of approaches from fuzzy logics to probabilistic networks [6, 15, 10, 7, 13, 4, 12, 11, 5, 1]. The difficulty lies in creating a knowledge representation with the right blend of ***flexibility*** and ***sound semantics***. For the human expert and knowledge engineer, flexibility and intuitiveness eases the acquisition and organization of knowledge for the target domain. On the other hand, sound and formal semantics prevents confusion concerning the interaction between the different sources of uncertainty. This serves to avoid anomalous reasoning behaviors and guarantees internal consistency.

Many other pragmatic concerns also drive our choice of knowledge representations such as accessibility and incompleteness. By accessibility, we mean whether the information required by the knowledge representation (in addition to that of the target domain) is actually available. For example, in a Bayesian network [7], certain conditional probabilities may not exist or are not meaningful in the target domain. Accessibility immediately leads to completeness. In the real world, having complete information is typically unattainable, thus a representation must have the ability to handle and recognize incompleteness as it occurs.

Most agree that encoding knowledge in terms of logical "if-then" style rules is the simplest and most intuitive approach to organization [3]. While uncertainty can be explicitly modeled as various exceptions to the rules, the number of such exception rules quickly explodes. On the other hand, probability theory is and has been an accepted language both for the description of uncertainty and for making inferences from incomplete knowledge. However, the general language of probabilities is too unconstrained making it a poor mechanism for easily organizing and manipulating information. Without additional knowledge such as independence conditions, the various sources of uncertainty can not be resolved or combined.

In this paper, we present a new model for knowledge representation under uncertainty called *Bayesian Knowledge-Bases (abbrev.* BKB*s)*. BKBs marry the strong formalisms for probability theory with an "if-then" rule structure. We will prove that BKBs provide a sound and consistent representation of knowledge and that the results they generate will not be inconsistent even in the face of incompleteness. Yet, even with all this formalism, BKBs still retain the flexibility of "if-then" style rules by providing intuitive probabilistic semantics.

Section 2 presents the basic foundations, models and properties necessary to define BKBs. Next, in Section 3, the semantics of BKBs are precisely defined and then proven to be consistent with probabilistic theory. Finally, in Section 4, we summarize our results and present our conclusions.

2 Foundations

We define a ***knowledge-base*** to consist of two components, a ***knowledge-graph*** and an *inferencing scheme*. The knowledge-graph represents objects/world states and the relationships between them. The inferencing scheme details how to manipulate or interpret the information in the knowledge-graph. The latter can be represented as subgraphs of the knowledge-graph.
NOTATION. $\Re$ denotes the real numbers, $\Re^+$ denotes the non-negative reals, and Φ denotes the empty set.

We define our knowledge-graph as follows:

DEFINITION 2.1. *A* correlation-graph $G = (I \cup S, E)$ *is a directed graph such that* $I \cap S = \Phi$ *and* $E \subseteq \{I \times S\} \cup \{S \times I\}$. *Furthermore, for all* $a \in S$, (a, b) *and* (a, b') *are in* E *if and only if* $b = b'$. $\{I \cup S\}$ *are the* nodes *of* G *and* E *are the* edges *of* G. *A node in* I *is called an* instantiation-node (abbrev. I-node) *and a node in* S *is called a* support-node (abbrev. S-node).

I-nodes represent the various states of the world such as the truth or falsity of a proposition. S-nodes, on the other hand, explicitly embody the relationships between the I-nodes. See Figure 2.1.
NOTATION. Let a be any node in $I \cup S$. $D_a^G = \{b|(b, a) \in E\}$ are the immediate predecessors of a in graph G.

Let π be a partition on I. Intuitive, π will denote the groups of I-nodes (states) which are mutually exclusive. For example, this can be used to represent random variables with discrete but multiple instantiations.
NOTATION. $|\cdot|$ denotes cardinality.

DEFINITION 2.2. G *is said to* I-respect π *if for all cells* σ *in* π, *for any S-node* $b \in S$, $|\sigma \cap (D_b^G \cup \{a\})| \leq 1$ *whenever* $(b, a) \in E$.

Basically, mutually exclusive I-nodes cannot be directly related to each other through the S-nodes. Next, we define mutual exclusion between S-nodes.

DEFINITION 2.3. *Two S-nodes* b_1 *and* b_2 *in* S *are said to be* mutually exclusive *with respect to* π *if there exists two distinct I-nodes* c_1 *and* c_2 *in some cell* σ *of* π *such that* $c_1 \in D_{b_1}^G$ *and* $c_2 \in D_{b_2}^G$.

DEFINITION 2.4. G *is said to* S-respect π *if for all I-nodes* a *in* I, *any two distinct S-nodes* b_1 *and* b_2 *in* D_a^G *are mutually exclusive.*

DEFINITION 2.5. G *is said to* respect π *if* G *both I-respects and S-respects* π.

To complete our knowledge-graph, we define a function w from S to $\Re$. This serves as the mechanism for handling uncertainty in the relationships.

DEFINITION 2.6. *A* knowledge-graph K *is a 3-tuple* (G, w, π) *where*

- $G = (I \cup S, E)$ *is a correlation-graph*
- w *is a function from* S *to* $\Re^+ \cup \{\infty\}$. *For each* $a \in S$, $w(a)$ *is the* weight *of* a.

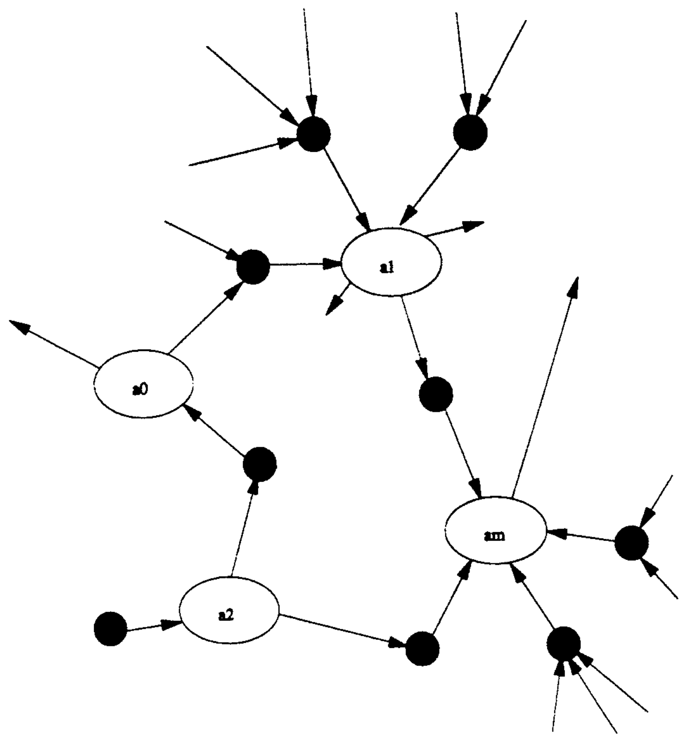

FIG. 2.1. Example correlation graph. The blackened nodes represent S-nodes. and the rest represent I-nodes. Note that each S-node has at most a single outgoing edge.

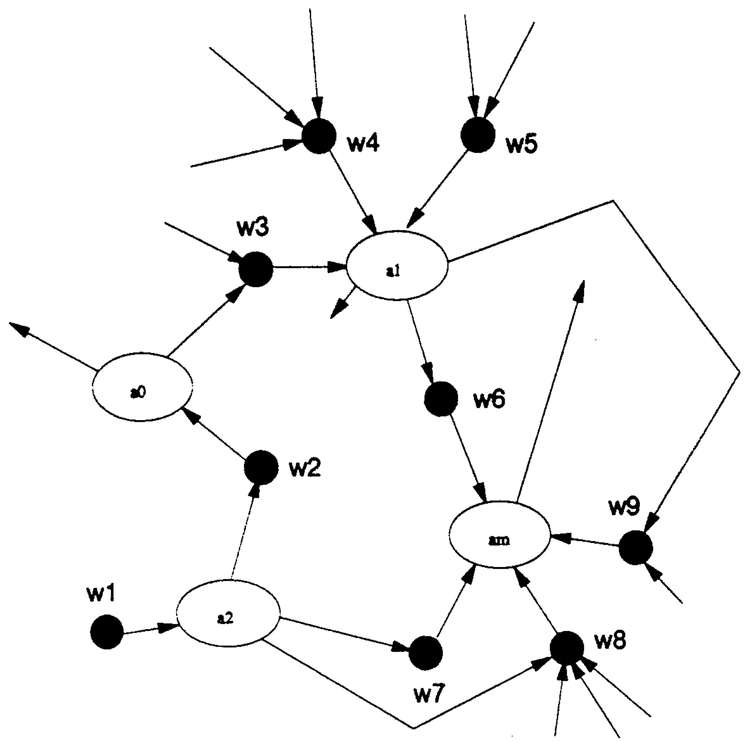

FIG. 2.2. Example of a knowledge graph. Let π consist of cells $\{a0\}$, $\{a1, a2\}$, and $\{am\}$.

- *π is a partition on I.*

and G respects π.

We now formally define our second component, the inferencing scheme. We begin by defining basic properties necessary to inferencing over our knowledge-graphs. Such properties include preserving the relationships between the objects in our graph, etc.

Let $r = (I' \cup S', E')$ be some subgraph of our correlation-graph $G = (I \cup S, E)$ where $I' \subseteq I$, $S' \subseteq S$, and $E' \subseteq E$. Furthermore, r has a *weight* $\omega(r)$ defined as follows:

$$\omega(r) = \sum_{s \in S'} w(s).$$

DEFINITION 2.7. *An I-node $a \in I'$ is said to be* well-supported *in r if there exists an edge (b, a) in E'. Furthermore, r is said to be* well-supported *if for all*

I-nodes a in I', a is well-supported.

Each I-node must have an incoming S-node in r.

DEFINITION 2.8. *An S-node $b \in S'$ is said to be* well-founded *in r if for all $(a,b) \in E$, $(a,b) \in E'$. Furthermore, r is said to be* well-founded *if for all S-nodes b in S', b is well-founded.*

If an S-node b is present in r, then all incoming I-nodes (conditions) to b in G must also be present in r.

DEFINITION 2.9. *An S-node $b \in S'$ is said to be* well-defined *in r if there exists an edge $(b,a) \in E'$. Furthermore, r is said to be* well-defined *if for all S-nodes b in S', b is well-defined.*

Each S-node in r must support some I-node in r.

DEFINITION 2.10. *r is said to be an* inference *over K if r is*

- *well-supported,*
- *well-founded,*
- *well-defined,*
- *acyclic, and*
- *For all cells σ in π, $|I' \cap \sigma| \leq 1$.*

Furthermore, r is said to be a complete inference *over K if*

- *For all cells σ in π, $|I' \cap \sigma| = 1$.*

Given the knowledge graph in Figure 2.2, one possible inference can be seen in Figure 2.3. An inference scheme for our knowledge-base will consist of inferences over our knowledge-graph.

DEFINITION 2.11. *A* knowledge-base *$\mathcal{K}$ is an ordered pair (K,R) where $K = (G,w,\pi)$ is a knowledge-graph and R is a non-empty collection of inferences over K called the* inference graphs.

With respect to our knowledge-base, we can derive properties for inference graphs which are especially pertinent to knowledge-bases.

THEOREM 2.1. *Let $r = (I' \cup S', E') \in R$. For each I-node $a \in I'$, there exists a unique S-node $b \in S'$ such that $(b,a) \in E'$.*

Proof. Since r is well-supported, there exists $b \in S'$ such that $(b,a) \in E'$. Now, assume $b' \in S'$ and $(b',a) \in E'$ but $b \neq b'$. Since G S-respects π, b and b' must be mutually exclusive. This implies there exists two distinct I-nodes c and c' in some cell σ in π such that $c \in D_b^G$ and $c' \in D_{b'}^G$. Thus, c and c' are both in I'. However, this violates Definition 2.10. Contradiction. □

Theorem 2.1 simply states that each I-node in an inference graph is supported by one and only one S-node.

PROPOSITION 2.2. $|I'| = |S'|$.

S-nodes not only serve as our mechanisms for representing uncertainty but also serves to "support" or "explain" the I-node it dominates. Since the I-nodes

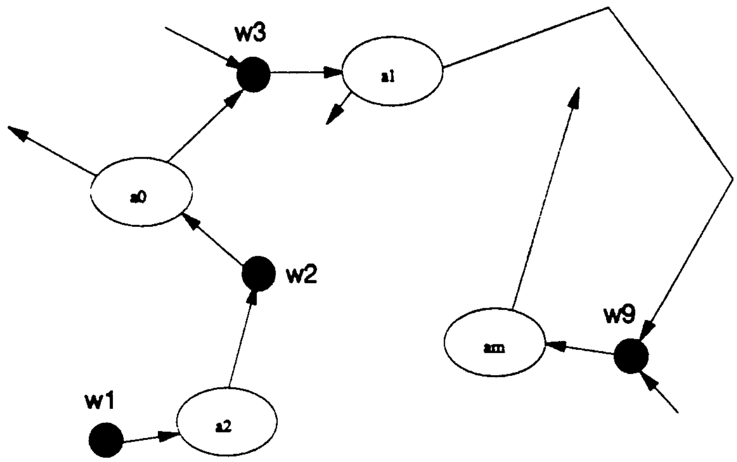

FIG. 2.3. Example of an inference. Let π consist of cells $\{a0\}$, $\{a1, a2\}$, and $\{am\}$.

are used to describe the various states of the world, ambiguity implies that there may be multiple explanations for a single state.

Let $r_1 = (I_1 \cup S_1, E_1)$ and $r_2 = (I_2 \cup S_2, E_2)$ be inference graphs over K.

THEOREM 2.3. *The following are equivalent:*

- $I_1 \subseteq I_2$.
- $S_1 \subseteq S_2$.
- $E_1 \subseteq E_2$.
- r_1 *is a subgraph of* r_2.

Proof.

Claim: $I_1 \subseteq I_2$ implies $S_1 \subseteq S_2$.

Assume $I_1 \subseteq I_2$ but S_1 is not a subset of S_2. This implies there exists some $b \in S_1$ but not in S_2. Since r_1 is well-defined, there exists some $a \in I_1$ such that $(b, a) \in E_1$. Thus, $a \in I_2$. Since r_2 is well-supported, there exists some $b' \in S_2$ such that $(b', a) \in E_2$. $b' \neq b$ and b must be mutually exclusive to b'. This implies there exists some c and c' in some cell σ of π such that $c \in D_b^G$ and $c' \in D_{b'}^G$. Since both r_1 and r_2 are well-founded, $c \in I_1$ and $c' \in I_2$. Thus, c and c' are both in I_2. Contradiction.

Claim: $S_1 \subseteq S_2$ implies $E_1 \subseteq E_2$.

Assume $S_1 \subseteq S_2$ but E_1 is not a subset of E_2. This implies there exists some $(a, b) \in E_1$ but not in E_2. If $a \in S_1$, then $a \in S_2$. By Definition 2.1, b is unique. Since r_2 is well-defined, $(a, b) \in E_2$. Thus, a must be in I_1. Since r_2 is well-founded and $b \in S_2$, (a, b) must also be in E_2. Contradiction.

Claim: $E_1 \subseteq E_2$ implies r_1 is a subgraph of r_2.

Assume $E_1 \subseteq E_2$ but r_1 is not a subgraph of r_2. This implies there exists some $a \in I_1 \cup S_1$ but not in $I_2 \cup S_2$. If $a \in I_1$, then since r_1 is well-supported, there exists some $b \in S_1$ such that $(b, a) \in E_1$. Thus, a must be in S_1. Since r_1 is well-defined, there exists some $b \in I_1$ such that $(a, b) \in E_1$. Contradiction.

Clearly, if r_1 is a subgraph of r_2, then $I_1 \subseteq I_2$. □

Theorem 2.3 provides another level of consistency in that inferences over portions of the shared subgraphs are identical.

COROLLARY 2.4. *The following are equivalent:*

- $I_1 = I_2$.
- $S_1 = S_2$.
- $E_1 = E_2$.
- $r_1 = r_2$.

Corollary 2.4 above says that each state of the world is uniquely determined by an inference graph in $\mathcal{K}$. Hence, $\mathcal{K}$ avoids ambiguity.

With respect to uncertainty, we have the following:

PROPOSITION 2.5. *If $S_1 = S_2$, then $\omega(r_1) = \omega(r_2)$.*

Key Theorem 2.1. *If $I_1 = I_2$, then $\omega(r_1) = \omega(r_2)$.*

Proof. This follows from Corollary 2.4 and Proposition 2.5. □

Key Theorem 2.1 lays the foundations for our Bayesian Knowledge-Base as we shall see later.

DEFINITION 2.12. *Let r_1 and r_2 be two subgraphs of G. r_1 is said to be* compatible *with r_2 if there does not exist two distinct I-nodes a_1 and a_2 in some cell σ in π such that $a_1 \in I_1$ and $a_2 \in I_2$. Otherwise, r_1 is said to be* incompatible *with r_2.*

This leads to the following properties on constructing inferences. Let $r = (I' \cup S', E') \in R$.

LEMMA 2.6. *Let $h = (I_h \cup S_h, E_h)$ be a subgraph of G such that r is a subgraph of h and for all $a \in I_h$, there exists at most one edge $(b, a) \in E_h$ for some $b \in S_h$. For each $a \in I_h - I'$, a cannot be an ancestor in h of any I-node $a' \in I'$.*

Proof. Assume there exists an $a \in I_h - I'$ such that a is an ancestor of some I-node $a' \in I'$. Without loss of generality, assume (a, b) and (b, a') are in E_h for some S-node $b \in S_h$. Since r is well-founded, this implies that if $b \in S'$, then $a \in I'$. Hence, b is not in S'. Since r is well-supported, there exists $b' \in S'$ such that $(b', a') \in E'$. Thus, both (b, a') and (b', a') are in E_h. Contradiction. □

From Lemma 2.6, we can conclude that inferences can be constructed in a bottom-up manner from other "smaller" subgraphs until we finish at a complete inference.

DEFINITION 2.13. *r_1 is said to be an* immediate parent *of r_2 if r_1 is a proper subgraph of r_2 and $|I_2 - I_1| = |S_2 - S_1| = 1$.*

Key Theorem 2.2. *If r_1 is a subgraph of r_2, then there exists a sequence of inferences $\{h_1, h_2, \ldots, h_n\}$ such that r_1 is the immediate parent of h_1, h_i is the immediate parent of h_{i+1} for i from 1 to $n-1$, and h_n is the immediate parent of r_2.*

Proof. From Definition 2.13 and Proposition 2.2, it follows that $n = |I_2 - I_1|$. Since r_1 is a subgraph of r_2, from Lemma 2.6, for any $a \in I_2 - I_1$, a cannot be an ancestor of any node in I_1. Since all inferences are acyclic, there must exist some I-node $a \in I_2 - I_1$ such that $D_b^{r_2} \subseteq I_1$ for some $(b, a) \in E_2$. We can construct $h_1 = (I_{h_1} \cup S_{h_1}, E_{h_1})$ from r_1 as follows: $I_{h_1} = I_1 \cup \{a\}$, $S_{h_1} = S_1 \cup \{b\}$, and $E_{h_1} = E_1 \cup \{(b, a)\} \bigcup_{c \in D_b^{r_1}} \{(c, b)\}$. We can easily show that h_1 is an inference over K and h_1 is a subgraph of r_2. We can construct the remaining h_is recursively. □

Key Theorem 2.2 further states that such a construction can occur as a sequence of inferences.

Definition 2.14. *An I-node $a \in I$ is said to be a root of K if there exists a S-node $b \in S$ such that $(b, a) \in E$ and $D_b^G = \Phi$. We denote the set of all roots of K by $V(K)$.*

For each root $a \in V(K)$, the corresponding root-subgraph *for a is consists solely of a, its associated S-node, and the edge between them.*

Clearly, all root-subgraphs are inferences over K.

Theorem 2.7. *If $a \in V(K)$, there exists a unique S-node $b \in S$ such that $(b, a) \in E$.*

Proof. This follows from the fact that $D_b^G = \Phi$ and that G S-respects π. □

Proposition 2.8. *All inferences contain one or more root-subgraphs.*

Combining Proposition 2.8 with Key Theorem 2.2, all possible inferences can be constructed from the root-subgraphs.

Theorem 2.9. *If r_1 is compatible with r_2, then the subgraph $r = (\{I_1 \cap I_2\} \cup \{S_1 \cap S_2\}, E_1 \cap E_2)$ is an inference over K.*

Proof. Clearly, r is acyclic and for all cells σ in π, $|(I_1 \cap I_2) \cap \sigma| \leq 1$. Furthermore, r can be easily shown to be well-founded and well-defined. Finally, based on Theorem 2.1, r can be shown to be well-supported. □

Given two compatible inferences, there is a common "core" between them which is also an inference and from which they may be constructed from.

Next, consider the following properties concerning the weights of the inferences.

Notation. Let σ be a cell in π. $\mathsf{span}_\sigma = \bigcup_{a \in \sigma} D_a^G$ denotes all the S-nodes which support the I-nodes in σ.

DEFINITION 2.15. K *is* locally normalized *if for each cell σ in π,*

$$\sum_{b \in S} e^{-w(b)} \leq 1$$

whenever $S \subseteq \mathsf{span}_\sigma$ and no two S-nodes in S are mutually exclusive.

Let $H = \{h_1, h_2, \ldots, h_n\}$ where $n > 1$ be a set of mutually incompatible inferences over K where $h_i = (I_i \cup S_i, E_i)$ for $i = 1, \ldots, n$. Let $C = \cap_{i=1}^{n} I_i$.

Key Theorem 2.3. *If K is locally normalized, then for any non-empty collection of mutually incompatible inferences H over K,*

$$\sum_{h \in H} e^{-\omega(h)} \leq 1.$$

Proof. We prove this by induction on the number of S-nodes in a given knowledge-graph. If a knowledge-graph has exactly one S-node and is locally normalized, then it is trivial to show that the knowledge-graph itself satisfies the above property.

Assume that the theorem is true for any locally normalized knowledge-graph of up to $n-1$ S-nodes.

Let K be a locally normalized knowledge-graph with exactly n S-nodes. Construct knowledge-graph $K^u = (G^u, w^u, \pi^u)$ where $G^u = (I^u \cup S^u, E^u)$, $I^u = \bigcup_{i=1}^{n} I_i$, $S^u = \bigcup_{i=1}^{n} S_i$, $E^u = \bigcup_{i=1}^{n} E_i$, $w^u(b) = w(b)$ for all $b \in S^u$, and $\pi^u = \{\sigma \cap I^u | \sigma \in \pi\}$. Clearly, K^u is locally normalized. If K^u is a proper subgraph of K, then K^u has less S-nodes than K and we are done.

Assume the $K^u \equiv K$. From Proposition 2.8, $V(K) \neq \Phi$. Let $a \in V(K)$. From Theorem 2.7, let b be the unique S-node in S such that $(b, a) \in E$.

Let $\sigma = \{a = a_0, a_1, a_2, \ldots, a_m\}$ be the cell in π containing a. If σ only contains a and since $a \in V(K)$, we can partition H into two sets H_a and $H - H_a$ where H_a are the inferences which contain a. Thus,

$$\begin{aligned}
\sum_{h \in H} e^{-\omega(h)} &= \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H_a} e^{\omega(h)} \\
&= \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H_a} e^{-\sum_{d \in S_h} w(d)} \\
&= \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H_a} e^{-w(b) - \sum_{d \in S_h - \{b\}} w(d)} \\
&= \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H_a} e^{-w(b)} e^{-\sum_{d \in S_h - \{b\}} w(d)} \\
&\leq \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H_a} e^{-\sum_{d \in S_h - \{b\}} w(d)}
\end{aligned}$$

where S_h are the S-nodes corresponding to h. Construct a new knowledge-graph K' from K be removing nodes a and b. Similarly, construct $H' = \{h'_1, \ldots, h'_n\}$ from H by eliminating a and b from the inferences. Clearly, the inference in $H - H_a$ are not affected. Let H'_a be the corresponding inferences found in H_a. Thus,

$$\begin{aligned}
\sum_{h \in H'} e^{-\omega(h)} &= \sum_{h \in H' - H'_a} e^{\omega(h)} + \sum_{h \in H'_a} e^{\omega(h)} \\
&= \sum_{h \in H' - H'_a} e^{\omega(h)} + \sum_{h \in H'_a} e^{-\sum_{d \in S'_h} w(d)} \\
&= \sum_{h \in H' - H'_a} e^{\omega(h)} + \sum_{h \in H' - H'_a} e^{-\sum_{d \in S_h - \{b\}} w(d)} \\
&= \sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H - H_a} e^{-\sum_{d \in S_h - \{b\}} w(d)}
\end{aligned}$$

where S'_h are the S-nodes corresponding to h'. Since K' has $n-1$ S-node,

$$\sum_{h \in H - H_a} e^{\omega(h)} + \sum_{h \in H - H_a} e^{-\sum_{d \in S_i - \{b\}} w(d)} \leq 1.$$

It follows then,

$$\sum_{i=1}^{n} e^{-\omega(h_i)} \leq 1.$$

Otherwise, if $m > 0$, then partition H into three disjoint subsets H_x, H_0, H_1 where

- $h_i \in H_x$ if and only if $I_i \cap \sigma = \Phi$.
- $h_i \in H_0$ if and only if $a_0 \in I_i$.
- $h_i \in H_1$ if and only if $\{a_1, \ldots, a_m\} \cap I_i \neq \Phi$.

Since $a_0 \in V(K)$, let b_0 be its unique associated S-node. Let B be the set of S-nodes found in H_1 that support the remaining I-nodes $\{a_1, \ldots, a_m\}$.

Now, construct a new K' from K as follows:

1. Introduce new I-node $\overline{a_0}$ and new S-node $\overline{b_0}$ for $\overline{a_0}$. Let $w(\overline{b_0}) = -\ln\left[1 - e^{-w(b_0)}\right]$.
2. Replace cell σ by $\sigma_1 = \{a_0, \overline{a_0}\}$ and $\sigma_2 = \{a_1, \ldots, a_m\}$.
3. For each $b \in B$, introduce $(\overline{a_0}, b)$ and change $w(b)$ to $w(b) - w(\overline{b_0})$.

From our construction, K' is also a locally normalized knowledge-base. See Figures 2.4 and 2.5 for an example transformation.

Next, construct H' from H as follows:

- For all $h \in H_x \cup H_0$, h is still a valid inference over K'.
- For all $h \in H_1$, introduce $\overline{a_0}$ and $\overline{b_0}$ as well as the additional arcs and changes in weights outlined in the construction of K' above.

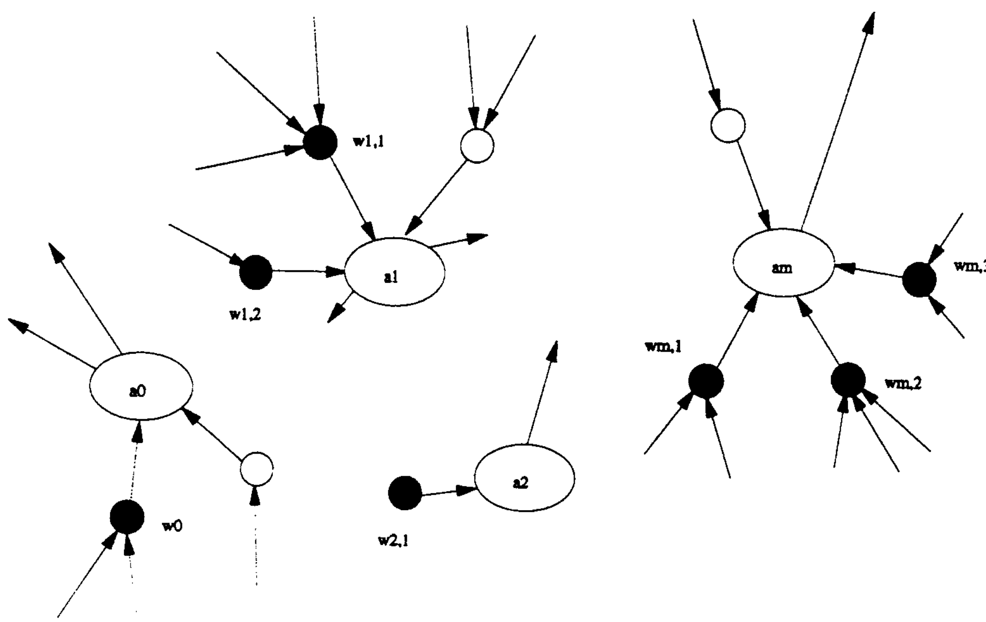

FIG. 2.4. Original knowledge-base where the darkened circular nodes represent S-nodes found in $B \cup \{b_0\}$ and wi, j represent their current weights.

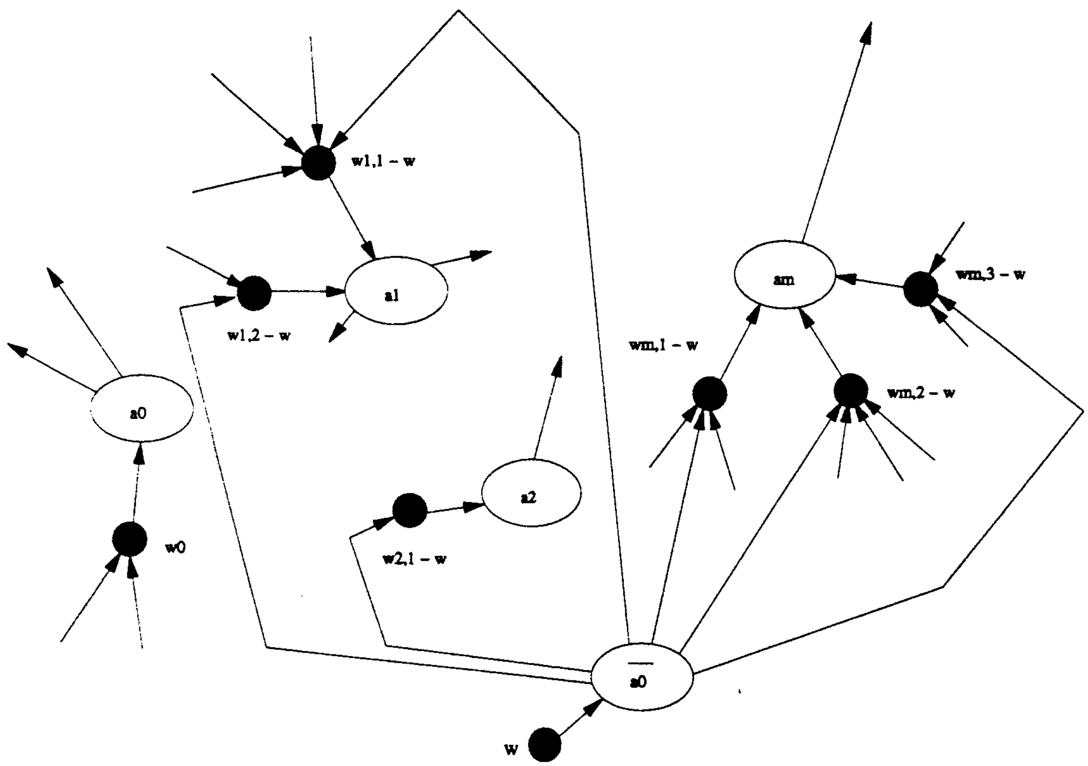

FIG. 2.5. Transformation to new knowledge-base where $w = w(\overline{b_0})$.

For each $h_i \in H$, let $h'_i \in H'$ be the corresponding inference and H'_x, H'_0, H'_1 be the corresponding partitions.

Clearly, $\omega(h_i) = \omega'(h'_i)$ and

$$\sum_{i=1}^{n} e^{-\omega(h_i)} = \sum_{i=1}^{n} e^{-\omega'(h'_i)}.$$

Since $w(b_0)$ is shared by all inferences in H'_0,

$$\begin{aligned}
&\sum_{h\in H'_0} e^{-\omega'(h)} \\
&\quad = \sum_{h\in H'_0} e^{-w'(b_0)} e^{-\omega'(h)+w'(b_0)} \\
&\quad = e^{-w'(b_0)} \sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)}.
\end{aligned}$$

Similarly, since $w'(\overline{b_0})$ is shared by all inferences in H'_1,

$$\begin{aligned}
&\sum_{h\in H'_1} e^{-\omega'(h)} \\
&\quad = \sum_{h\in H'_1} e^{-w'(\overline{b_0})} e^{-\omega'(h)+w'(\overline{b_0})} \\
&\quad = e^{-w'(\overline{b_0})} \sum_{h\in H'_1} e^{-\omega'(h)+w'(\overline{b_0})}.
\end{aligned}$$

- **Case 1.** $\sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)} \geq \sum_{h\in H'_1} e^{-\omega'(h)+w'(\overline{b_0})}$.

$$\begin{aligned}
&\sum_{h\in H'_x} e^{-\omega'(h)} + e^{-w'(b_0)} \sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)} + e^{-w'(\overline{b_0})} \sum_{h\in H'_1} e^{-\omega'(h)+w'(\overline{b_0})} \\
&\leq \sum_{h\in H'_x} e^{-\omega'(h)} + e^{-w'(b_0)} \sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)} + e^{-w'(\overline{b_0})} \sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)} \\
&= \sum_{h\in H'_x} e^{-\omega'(h)} + \sum_{h\in H'_0} e^{-\omega'(h)+w'(b_0)}
\end{aligned}$$

since $e^{-w'(b_0)} + e^{-w'(\overline{b_0})} = 1$.

Construct a new knowledge-base K'' from K' by eliminating all I-nodes $\{\overline{a_0}, a_1, a_2, \ldots, a_m\}$ and their attendant support nodes from K'. Furthermore, remove any S-nodes which are immediate descendants of the I-nodes. If a situation occurs where an I-node has no support nodes, also eliminate that I-node. Repeat until no I-nodes need to be removed.

Next, let $w''(b_0) = 0$. The inferences in H'_x and H'_0 do not need to be changed. In fact, for any $h \in H'_0$, its new weight under K'' is precisely $\omega''(h) = \omega'(h) - w'(b_0)$. K'' is also locally normalized.
By eliminating the I-nodes and S-nodes, K'' has fewer than n S-nodes. By induction,

$$\sum_{h \in H'_x} e^{-\omega''(h)} + \sum_{h \in H'_0} e^{-\omega''(h)} \leq 1.$$

Thus,

$$\sum_{h \in H} e^{-\omega(h)} \leq 1.$$

- **Case 2.** $\sum_{h \in H'_0} e^{-\omega'(h)+w'(b_0)} \leq \sum_{h \in H'_1} e^{-\omega'(h)+w'(\overline{b_0})}$.
Similar to Case 1, we construct a new knowledge-base K'' by eliminating I-nodes $\{a_0, \overline{a_0}\}$. Again, K'' is locally normalized and has fewer than $n-1$ S-nodes.

□

Key Theorem 2.3 states that local normalization implies global normalization. This property will play a key role in our semantics for these knowledge-bases. In addition, this leads to the following even stronger result concerning local normalization.

Key Theorem 2.4. *Let H be a set of mutually incompatible inferences over K such that for each h in H, h is a supergraph of r. If $\mathcal{K}$ is locally normalized, then*

$$\sum_{h \in H} e^{-\omega(h)} \leq e^{-\omega(r)}.$$

Proof. Let $H = \{h_1, h_2, \ldots, h_n\}$ and C' be the subset of $\cup_{i=1}^{n} I_i$ such that if $a \in C'$, then $(\sigma_a - \{a\}) \cap I_i = \Phi$ for $i = 1, \ldots, n$ where σ_a is the cell in π containing a. For each h_i, $\omega(h_i) = \omega(r) + \beta_i$ for some $\beta_i \geq 0$ since r is a subgraph of h_i and from Theorem 2.9.

From Key Theorem 2.3,

$$\sum_{i=1}^{n} e^{-\omega(h_i)} \leq 1.$$

Thus, it follows that

$$e^{-\omega(r)} \sum_{i=1}^{n} e^{-\omega(\beta_i)} \leq 1.$$

Construct a new knowledge-base $\mathcal{K}'$ from $\mathcal{K}$ by eliminating all I-nodes from K not in $\cup_{i=1}^{n} I_i$ as well as any S-nodes which are immediate parents or immediate children of these I-nodes. Furthermore, eliminate all S-nodes which do not occur in $\cup_{i=1}^{n} S_i$. For all remaining S-nodes which support the I-nodes in C', change

their weights to 0. Clearly, $\mathcal{K}'$ is a locally normalized knowledge-base. From Key Theorem 2.3,

$$\sum_{i=1}^{n} e^{-\omega'(h_i)} \leq 1$$

and consequently

$$\sum_{i=1}^{n} e^{-\beta_i} \leq 1.$$

Therefore,

$$\sum_{h \in H} e^{-\omega(h)} \leq e^{-\omega(r)}.$$

□

3 Semantics

In this section, we describe how knowledge-bases maybe used to represent probabilistic information.

Let $\mathcal{A} = \{A_1, A_2, \ldots, A_n\}$ be a set of discrete random variables[2] and let $R(A_i) = \{a_{i,1}, a_{i,2}, \ldots, a_{i,m_i}\}$ represent the set of possible instantiations to r.v. A_i for $i = 1, \ldots, n$.

DEFINITION 3.1. *An* instantiation *is an ordered pair* (A, a) *where* $A \in \mathcal{A}$ *and* $a \in R(A)$. *(An instantiation* (A, a) *is also denoted by* $A = a$.*) A collection of instantiations* X *is called an* instantiation-set *iff* $(A, a), (A, a')$ *in* X *implies* $a = a'$.

An instantiation represents the event when a r.v. takes on a value from its range. Given an instantiation-set, we can define the notion of the *span* of an instantiation-set.

DEFINITION 3.2. *Given an instantiation-set* X*, we define the* span *of* X, $\mathsf{span}(X)$, *to be the collection of r.v.s in the first coordinate of the instantiations. Furthermore, an instantiation-set* X *is said to be* complete *if and only if* $\mathsf{span}(X) = \mathcal{A}$.

The span of an instantiation-set simply denotes the r.v.s which have been instantiated.

Let $\mathcal{K} = (K, R)$ with $K = (G, w, \pi)$ be a locally normalized knowledge-base where π consists of cells $\{\sigma_1, \sigma_2, \ldots, \sigma_n\}$ such that there exists a one-to-one and onto mapping functions $f_i : R'(A_i) \rightarrow \sigma_i$ and $R'(A_i) \subseteq R(A_i)$ for $i = 1, \ldots, n$.

Let $r = (I' \cup S', E')$ be an inference over K. According to Definition 2.10, there exists at most one I-node from each σ_i in r. Hence, we can construct an instantiation set X_r from r as follows: $(A_i, a_{i,j}) \in X_r$ if and only if $f_i(a_{i,j}) \in I'$. Intuitively, if I-node $f_i(a_{i,j}) \in I'$, this implies that r.v. $A_i = a_{i,j}$.

[2] Random variables will be abbreviated r.v.

PROPOSITION 3.1. *r is complete if and only if X_r is complete.*

THEOREM 3.2. *For each instantiation-set X, there corresponds at most one inference r_X over K.*

Proof. This follows from Theorem 2.3 and Corollary 2.4. $\Box$

From Theorem 3.2 and Key Theorem 2.1, we can associate a unique weight $\omega(r)$ to each X_r. Our goal now is to interpret this weights probabilistically.

Since each $\omega(r)$ is guaranteed to be non-negative, $0 \leq e^{-\omega(r)} \leq 1$. In addition, from Key Theorem 2.3, given any set of mutually incompatible inference $\mathcal{H} = \{h_1, h_2, \ldots, h_l\}$,

$$\sum_{i=1}^{l} e^{-\omega(h_i)} \leq 1.$$

Thus, given an inference r and the corresponding instantiation-set X_r, we can interpret $e^{-\omega(r)}$ as the joint probability for X_r.

Now, consider the S-nodes and their weights w. Let $b \in S$ and $(b, b_1) \in E$. Note, b_1 is unique to b from Definition 2.1.

Since G S-respects π, we can associate a unique conditional probability with each S-node. Let $D_b^G = \{b_2, b_3, \ldots, b_k\}$. Without loss of generality, assume $b_i \in \sigma_i$ for $i = 1, \ldots, k$. For b, associate the conditional probability

$$P(A_1 = a_{1,j_1} | A_2 = a_{2,j_2}, \ldots, A_k = a_{k,j_k})$$

where $b_i = f_i(a_{i,j_i})$ for $i = 1, \ldots, k$. Furthermore, we assume the following conditional independence for the conditional probabilities associated with each S-node:

$$\begin{aligned} P(A_1 = a_{1,j_1} | A_2 = a_{2,j_2}, \ldots, A_k = a_{k,j_k}, C) = \\ P(A_1 = a_{1,j_1} | A_2 = a_{2,j_2}, \ldots, A_k = a_{k,j_k}) \end{aligned} \tag{1}$$

for any instantiation-set C such that $\mathsf{span}(C) \cap \{A_1, \ldots, A_k\} = \Phi$.

Returning to inference r, without loss of generality, assume $I' = \{c_1, c_2, \ldots, c_m\}$ and $c_i \in \sigma_i$ for $i = 1, \ldots, m$. The joint probability $P(X_r)$ is

$$P(A_1 = a_{1,j_1} | A_2 = a_{2,j_2}, \ldots, A_k = a_{m,j_m})$$

where $c_i = f_i(a_{i,j_i})$ for $i = 1, \ldots, m$. From Definition 2.7 and Theorem 2.1, each c_i is supported by a single S-node, say d_i.

Since r is acyclic, this induces a topological ordering on the I-nodes in I'. Without loss of generality, assume $\{c_1, \ldots, c_m\}$ is the ordering and c_i is not a descendant of c_j for $i > j$. Through the chain rule, we can rewrite $P(X_r)$ as

$$P(X_r) = \prod_{i=1}^{m} P(A_i = a_{i,j_i} | A_{i+1} = a_{i+1,j_{i+1}}, \ldots, A_m = a_{m,j_m}). \tag{2}$$

According to our S-nodes d_i and our conditional independence (1), we can replace each term in (2) by the corresponding conditional probability, i.e., for the term with A_i as the r.v. being conditioned upon, we can replace it with the conditional probability associated with S-node d_i. Thus, our joint probability $P(X_r)$ can be computed as a product of the conditionals.

Finally, if we set the conditional probability for d_i to $e^{-w(d_i)}$, we have

$$\begin{aligned}\prod_{i=1}^{m} e^{-w(d_i)} &= e^{-\sum_{i=1}^{m} w(d_i)} \\ &= e^{-\omega(r)}.\end{aligned}$$

Applying these probabilistic semantics to $\mathcal{K}$ is the basis of Bayesian knowledge-bases (BKBs). Reasoning over BKBs can thus involve search for the most-probable inference in R. Carefully selecting the inferences which make up R provides a wide variety of different reasoning schemes. For example, to determine the most probable state of the world when given certain evidence can be achieved as follows: Let E be the instantiation-set representing the evidence. Our goal is to determine the complete instantiation-set X^* which maximizes

$$P(X^*|E) = \max_{X} P(X|E)$$

over all complete instantiation-sets. Note that

$$P(X|E) = \frac{P(X,E)}{P(E)} = \frac{P(X)}{P(E)}.$$

Since $P(E)$ is a constant, we only need to determine the most probable complete instantiation-set X^* such that $E \subseteq X^*$. Thus, we can accomplish this letting R be the set of all complete inferences that also contain the I-nodes associated with E.

Performing these reasoning computations can be accomplished in a wide variety of ways. This includes basic A* search and integer linear programming [8] for exact solutions or a genetic algorithms approach [2].

4 Conclusion

By carefully defining the mathematical model and formal semantics for Bayesian knowledge-bases, we were able to develop a new framework for knowledge representation which unifies the flexible and intuitive structure of "if-then" style rules with probability theory. The new knowledge representation is fully probabilistic yet also retains the full power of "if-then" style. Furthermore, BKBs are naturally adept at handling incompleteness and are especially well-suited for knowledge acquisition. (For more details on using BKBs for knowledge acquisition, see [9].)

Clearly, BKBs can subsume Bayesian networks [7] using a similar probabilistic semantics. Unlike Bayesian networks, BKBs can also handle various forms of cyclical knowledge. Finally, it seems likely that temporal information could also be unified into BKBs without necessarily resorting to time-slice approaches. Instead, time can be used as a variable to directly influence the level of uncertainty [14].

References

[1] Fahiem Bacchus. *Representing and Reasoning with Probabilistic Knowledge: A Logical Approach to Probabilities.* The MIT Press, 1990.

[2] Brett J. Borghetti, Edward M. Williams, and Eugene Santos, Jr. Inferencing over incomplete solution spaces with genetic algorithms for probabilistic reasoning. In *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Conference*, 1996.

[3] Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems.* Addison Wesley, 1984.

[4] A. P. Dempster. A generalization of bayesian inference. *J. Royal Statistical Society*, 30:205–47, 1968.

[5] David Heckerman. *Probabilistic Similarity Networks.* The MIT Press, 1991.

[6] N. J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1986.

[7] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[8] Eugene Santos, Jr. A linear constraint satisfaction approach to cost-based abduction. *Artificial Intelligence*, 65(1):1–28, 1994.

[9] Eugene Santos, Jr. and Darwyn O. Banks. Acquiring consistent knowledge. Technical Report AFIT/EN/TR96-01, Department of Electrical and Computer Engineering, Air Force Institute of Technology, 1996.

[10] Eugene S. Santos and Eugene Santos, Jr. Reasoning with uncertainty in a knowledge-based system. In *Proceedings of the International Symposium on Multiple-Valued Logic*, pages 75–81, 1987.

[11] G. A. Shafer. *Mathematical Theory of Evidence.* Princeton University Press, 1979.

[12] E. H. Shortliffe and B. G. Buchanan. A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23:351–379, 1975.

[13] Paul Thagard. Explanatory coherence. *Behavioral and Brain Sciences*, 12:435–502, 1989.

[14] Joel D. Young and Eugene Santos, Jr. Introduction to temporal bayesian networks. In *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Conference*, 1996.

[15] L. A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems*, 11:199–227, 1983.